Adaptive interventions for social welfare maximization in network games

Mehran Shakarami, Ashish Cherukuri, Nima Monshizadeh

Abstract— **We consider the problem of steering the actions of noncooperative players in quadratic network games to the social optimum. To this end, a central regulator modifies the marginal returns of the players, while the players' strategies are determined by continuous pseudo-gradient dynamics. Depending on the available information on the players parameters and network quantities, suitable static and dynamic intervention protocols are devised that maximize the social welfare at steady-state. We show that adaptive interventions can compensate for the lack of knowledge on network topology and coupling weights. Numerical examples are provided to demonstrate the effectiveness of the proposed interventions.**

I. Introduction

Network games are a standard tool for modeling and studying the interaction between a population of decision makers, i.e., players, whose individual payoffs depend on their own decisions/actions as well as the actions of their neighboring players, determined by an underlying interaction network. This class of games has appeared in a broad spectrum of applications such as studying crime networks [1], pricing in social networks [2], [3], public good provision [4], [5], firm competition [6], and telecommunication [7]. Influencing the outcomes of network games through interventions is extensively studied in economics, and various works have investigated the impact of the network topology on optimal policies, see e.g., [1], [8], [9].

From another perspective, self-interested/selfish behavior of the players in a noncooperative game entails degradation of performance in comparison to the scenarios where the players would cooperate to optimize the social welfare. Such deterioration in performance or loss of efficiency has lead to the definition of the price of anarchy [10], and its quantification is extensively studied in different settings such as resource allocation [11], congestion games [12], [13], and supply chains [14].

An active line of research concerns decreasing the price of anarchy and realigning the preferences of the players with the social optimum through interventions. In this setup, a central regulator provides incentives in order to coordinate the players and steer their strategies towards the social optimum. The main challenge that arises in this research topic is that suitable incentives depend on private information of the agents, generally unknown to the regulator [15]. The celebrated Vickrey–Clarke–Groves (VCG) mechanism [16] is adopted in different disciplines, and especially in economics, to address this problem. In this setting, the mechanism generates a payment rule with the aim of incentivizing the players to report their private information to the regulator. This information is then used to reach to the social optimum, see [17] for more details on the topic.

Another methodology for reaching to the social optimum in noncooperative games involves exploiting control-theoretic tools. In this manner, the players do not report their private information, but their actions are observed over time by the regulator. The problem is then regarded as a feedback control problem where the desired outcome is the social optimum and the control effort is implemented through interventions [18]. Devising suitable control laws is straightforward when the regulator has *perfect information* on the game and the payoffs of the players, it becomes much more intricate when some of the players' private information and/or network level parameters are unknown. To overcome this lack of information, dynamical protocols are proposed in [18]–[20]. In [18], a dynamic pricing mechanism is devised that solves the problem for players with separable utility functions. When the utility functions are non-separable, side information is used in [19] for convergence to the social optimum. In particular, the pricing mechanism employs the utility functions evaluated at the social optimum. In the context of congestion control, the mechanism presented in [20] guarantees convergence assuming that the network manager knows the aggregate flow on each link as well as the delay-cost experienced by the users.

In this work, we present intervention protocols that can steer the outcome of quadratic network games to the solution of the social welfare maximization problem. The players are selfish and merely interested in maximizing their individual payoff functions. They do this by choosing their actions using continuous pseudo-gradient dynamics. The regulator, on the other hand, attempts to steer the players towards the social optimum by devising suitable interventions which modify the marginal returns of the players. We investigate multiple scenarios resulting from the limited access of the regulator to game information and the network structure. For each scenario, we propose a suitable intervention that is able to steer the players towards the desired social optimum. We analytically prove convergence of these intervention protocols, and complement our theoretical findings by numerical examples.

The structure of the paper is organized as follows: The problem formulation is given in Section II. Section III presents the intervention protocols that steer the actions of the players to the social optimum, and provides convergence

Mehran Shakarami, Ashish Cherukuri, and Nima Monshizadeh are with the Engineering and Technology Institute, University of Groningen, 9747AG, The Netherlands, m.shakarami@rug.nl, a.k.cherukuri@rug.nl, n.monshizadeh@rug.nl.

guarantees. Numerical examples are provided in Section IV. Concluding remarks and future research directions are stated in Section V.

Notation: The set of real numbers is denoted by $\mathbb{R}$. We denote the standard Euclidean norm by $\|\cdot\|$. The symbol $\mathbf{0}$ denotes a vector/matrix of all zeros. Given a vector $x \in \mathbb{R}^n$, we denote its i-th element of by $(x)_i$. For given vectors $x_1, \cdots, x_m \in \mathbb{R}^n$, we use the shorthand notation $\mathrm{col}(x_i) = \left[x_1^\top, \cdots, x_m^\top\right]^\top$. We write $P \succ \mathbf{0} (\prec \mathbf{0})$ to denote that the matrix $P = P^\top \in \mathbb{R}^{n\times n}$ is positive definite (negative definite). Given a matrix $P = P^\top \in \mathbb{R}^{n\times n}$, we denote its Frobenius norm by $\|P\|_{\mathrm{F}} = \sqrt{\mathrm{Tr}(P^\top P)}$ where $\mathrm{Tr}(\cdot)$ is the trace operator. Moreover, the notation $\lambda_i(P)$ with $i \in \{1, \ldots, n\}$ denotes the eigenvalues of P.

II. Problem formulation

We consider a game with the population of $\mathcal{I} := \{1, \ldots, n\}$ players/agents that interact repeatedly with a central regulator as well as with each other according to an underlying interaction network. We indicate the adjacency matrix of this network by $P \in \mathbb{R}^{n\times n}$ where $P_{ij} \in [0, 1]$ denotes the influence of player j's strategy/action on the utility function of player i. We assume that the network has no self loop, thus $P_{ii} = 0$ for all $i \in \mathcal{I}$, and the set of neighbors of player i is denoted by $\mathcal{N}_i = \{j \in \mathcal{I} \mid P_{ij} > 0\}$. The interaction network is undirected if $P_{ij} = P_{ji}$ for all $i, j \in \mathcal{I}$, otherwise it is directed.

Each player i is associated with a payoff function $U_i(x_i, z_i(x), u_i)$ that depends on her own action $x_i \in \mathbb{R}$, the aggregate of her neighbors' actions

$$z_i(x) := \sum_{j\in\mathcal{N}_i} P_{ij} x_j \tag{1}$$

with $x = \mathrm{col}(x_i)$, and an *intervention* $u_i \in \mathbb{R}$ which will be determined by the central regulator. We focus our attention on linear quadratic payoff functions of the form

$$U_i\big(x_i, z_i(x), u_i\big) = W_i\big(x_i, z_i(x)\big) + x_i u_i \tag{2}$$

with

$$W_i\big(x_i, z_i(x)\big) := -\frac{1}{2}x_i^2 + x_i\big(a z_i(x) + b_i\big), \tag{3}$$

where $a \in \mathbb{R}$ captures the impact of neighbors' actions and $b_i \in \mathbb{R}$ is the standalone marginal return. The payoff function W_i is used in the literature to model peer effects in social and economic settings, see e.g. [1], [21], [22]. In addition, the term $x_i u_i$ is included in (2) to capture the intervention of the central regulator in modifying the standalone marginal return b_i to $b_i + u_i$ [8].

The players are noncooperative and merely interested in maximizing their individual payoff functions by choosing their actions. This selfish behavior causes loss of efficiency with respect to the situation in which the players would cooperate to maximize the total payoff. The central regulator, on the other hand, is aimed at coordinating the players and avoiding the efficiency loss. To this end, she modifies the standalone marginal returns of the players through interventions.

In the next two subsections, we introduce the decision making process for the players and the regulator.

A. Players' strategy

Each player aims to maximize her individual payoff function U_i given the aggregated actions of her neighbors $z_i(x)$ and the current value of the intervention signal u_i. To capture this, we consider that the actions of the players evolve over time according to the following pseudo-gradient dynamics[1]:

$$\dot{x}_i(t) = \frac{\partial U_i}{\partial x_i}\big(x_i(t), z_i(x(t)), u_i(t)\big), \qquad \forall i \in \mathcal{I},$$

where $u_i(\cdot)$ is the intervention designed by the regulator. Noting the definition of $z_i(x)$ given by (1) and the fact that $P_{ii} = 0$, we can rewrite dynamics above as

$$\dot{x}_i(t) = -x_i(t) + a\sum_{j\in\mathcal{I}} P_{ij}x_j(t) + b_i + u_i(t). \tag{4}$$

In the case of no intervention, i.e., $u_i(\cdot) = 0$, the equilibrium of (4) coincides with the Nash equilibrium of the game[2], namely the action profile $\bar{x}$ satisfying

$$\bar{x}_i \in \arg\max_{y\in\mathbb{R}} U_i\big(y, z_i(\bar{x}), 0\big), \quad \forall i \in \mathcal{I}.$$

This can be written more explicitly as

$$(I - aP)\bar{x} = b. \tag{5}$$

Next we look at the problem from the regulator's side.

B. Regulator's objective

The central regulator aims to implement suitable interventions to coordinate the players and maximize the total payoff. More precisely, she aims at designing the intervention signal $u_i(t)$ such that the actions of the players converge to a *social optimum* x_{opt}, defined as a solution of the *social welfare* maximization problem:

$$x_{\text{opt}} \in \arg\max_{y\in\mathbb{R}^n} \sum_{i\in\mathcal{I}} W_i\big(y_i, z_i(y)\big), \tag{6}$$

where $y = \mathrm{col}(y_i)$ and W_i is given by (3). The necessary and sufficient condition for existence of a unique social optimum x_{opt} is given below:

Lemma II.1. *The social welfare maximization problem* (6) *has a unique solution if and only if*

$$\max_{i\in\mathcal{I}} a\,\lambda_i(P + P^\top) < 1. \tag{7}$$

Proof. If part: The inequality (7) implies that $-I + a(P + P^\top) \prec \mathbf{0}$, and thus the map $x \mapsto \sum_{i\in\mathcal{I}} W_i\big(x_i, z_i(x)\big)$ is strictly concave and so admits a unique maximizer [26, Prop. 1.1.2].

[1]See [23], [24] for further applications of continuous pseudo-gradient dynamics in the context of distributed Nash equilibrium seeking for noncooperative games.

[2]Existence of a Nash equilibrium readily follows from [25, Cor. 4.2].

Only if part: Suppose (7) does not hold. Then, the maximum eigenvalue of the symmetric matrix $-I+a(P+P^\top)$, denoted by μ, is nonnegative. Let $v \in \mathbb{R}^n$ be a corresponding eigenvector, and consider the social welfare function along the direction of v, namely the map

$$\alpha \mapsto \sum_{i\in\mathcal{I}} W_i\big((\alpha v)_i, z_i(\alpha v)\big) = \frac{1}{2}\mu\alpha^2\|v\|^2 + \alpha v^\top b,$$

where $(\alpha v)_i$ is the ith element of αv. Clearly, if $\mu > 0$, then the function grows unbounded as α increases to infinity and thus, the optimizer does not exist. When $\mu = 0$, then the matrix $-I + a(P + P^\top)$ is not invertible. This implies that the gradient of the social welfare function vanishes at infinitely many points, and thus (6) does not admit a unique maximizer. ☐

Motivated by Lemma II.1, we impose the following standing assumption throughout the paper.

Assumption II.2. The adjacency matrix $P \in \mathbb{R}^{n\times n}$ and the parameter $a \in \mathbb{R}$ satisfy $\max_{i\in\mathcal{I}} a\,\lambda_i(P + P^\top) < 1$. •

Remark II.3. The matrix $P + P^\top$ is symmetric with the diagonal elements equal to zero. This implies that the matrix $P + P^\top$ has only real eigenvalues, and its trace is zero. Hence,

$$\lambda_{\min}(P + P^\top) < 0 < \lambda_{\max}(P + P^\top).$$

It follows from the above inequalities that Assumption II.2 is satisfied if and only if either (i) $a > 0$ and $a\,\lambda_{\max}(P+P^\top) < 1$ or (ii) $a < 0$ and $a\,\lambda_{\min}(P + P^\top) < 1$. •

As a consequence of Assumption II.2, the social welfare function on the right hand side of (6) is strictly concave and thus admits a unique maximizer given by [26, Prop. 1.1.2]

$$x_{\text{opt}} = \big(I - a(P + P^\top)\big)^{-1} b. \tag{8}$$

Notice that this is different from the Nash equilibrium given in (5). Next, we aim at designing intervention mechanisms that asymptotically steer the players from the "selfish" behavior in (5) to the one in (8) which maximizes the social welfare.

III. Intervention protocols

First, we recall that by (4) the action profile evolves according to the following pseudo-gradient dynamics:

$$\dot{x}(t) = (-I + aP)x(t) + b + u(t). \tag{9}$$

Note that the matrix $-I + aP$ is Hurwitz since we have $-2I + a(P + P^\top) \prec \mathbf{0}$ (cf. Assumption II.2). Therefore, under the assumption that the regulator has full access to the game information, namely the pair (aP, b), a simple static *open-loop* intervention $u(t) \equiv (I - aP)x_{\text{opt}} - b$ suffices for convergence to x_{opt}. In the sequel, we look into the scenarios where such perfect information is not available to the regulator, hence more elaborate interventions are required.

A. Static feedback intervention

When the regulator has complete knowledge about the network and the impact of the players on each other, i.e., aP, a static state feedback intervention can be adopted to ensure convergence to the social optimum. This is formalized in the following proposition.

Proposition III.1. *Consider the pseudo-gradient dynamics in* (9). *Then, under the* static feedback *intervention $u(t) = aP^\top x(t)$, the action profile $x(t)$ exponentially converges to the social optimum x_{opt} given by* (8).

Proof. The proof directly follows from the expression of x_{opt} in (8) and the fact that the matrix $-I+a(P+P^\top)$ is negative definite (cf. Assumption II.2). ☐

B. Dynamic intervention with estimated social optimum

Next we consider the scenario where the regulator is not aware of the network information aP but has a reliable estimate of the social optimum x_{opt}. In this case, the regulator can resort to an integral control-based intervention

$$\dot{u}(t) = -\big(x(t) - x_s\big), \tag{10}$$

where $x_s \in \mathbb{R}^n$ is an estimation of the social optimum. Note that the above mechanism does not require any knowledge on the game parameters aP and b. We provide convergence guarantees for this integral control in the following proposition:

Proposition III.2. *Consider the pseudo-gradient dynamics* (9). *Then, under the intervention* (10)*, the action profile $x(t)$ exponentially converges to x_s.*

Proof. Let $u^* := (I - aP)x_s - b$, and consider the change of coordinates $(x, u) \mapsto (\tilde{x}, \tilde{u})$ with $\tilde{x} = x - x_s$ and $\tilde{u} = u - u^*$. In these coordinates, the overall closed-loop dynamics, consisting of (9) and (10), takes the form $\text{col}(\dot{\tilde{x}}, \dot{\tilde{u}}) = A\,\text{col}(\tilde{x}, \tilde{u})$ where

$$A = \begin{bmatrix} -I + aP & I \\ -I & \mathbf{0} \end{bmatrix}.$$

We prove stability by finding a matrix $M = M^\top \succ \mathbf{0}$ such that the Lyapunov inequality $A^\top M + MA \prec \mathbf{0}$ holds. Given some $\kappa > 0$, we define

$$M := \begin{bmatrix} I & -\kappa I \\ -\kappa I & I \end{bmatrix}.$$

Note that $M \succ \mathbf{0}$ for any $\kappa \in (0, 1)$. In addition, we have the following:

$$A^\top M + MA = \begin{bmatrix} -2(1-\kappa)I + a(P + P^\top) & \kappa(I - aP^\top) \\ \kappa(I - aP) & -2\kappa I \end{bmatrix}.$$

We use the Schur complement to deduce that the above matrix is negative definite if and only if

$$\kappa > 0,\ -2(1-\kappa)I + a(P+P^\top) + \frac{1}{2}\kappa(I - aP^\top)(I - aP) \prec \mathbf{0}.$$

It then follows from $-I + a(P + P^\top) \prec \mathbf{0}$ (cf. Assumption II.2) that there exists a sufficiently small $\kappa > 0$ such that the above inequality holds. As a result, the Lyapunov inequality is satisfied, and the dynamics $\operatorname{col}(\dot{\tilde{x}}, \dot{\tilde{u}}) = A\operatorname{col}(\tilde{x}, \tilde{u})$ is exponentially stable. This means that the solution $(x(t), u(t))$ converges to (x_s, u^*) exponentially fast, which concludes the proof. □

C. Adaptive intervention with known standalone marginal returns

Recall that in case the regulator knows aP or the social optimum x_{opt}, she could steer the players to the social optimum by implementing the previously discussed interventions. Here, we shift our focus to the case where both aP and x_{opt} are unknown to the regulator, and she merely has knowledge about the individual standalone marginal returns of the players b. It turns out that this substantially complicates the problem faced by the regulator. We restrict our attention in this subsection to the case of undirected networks, i.e. $P = P^\top$.

A natural approach to tackle this problem is to resort to adaptive control techniques which potentially allow to compensate for lack of complete knowledge on the system dynamics. However, there are certain obstacles that hinder an application of standard adaptive control schemes. First, a control design based on the regulation error $x(t) - x_{\text{opt}}$ is not feasible since x_{opt} is unknown. A second attempt would be to try to estimate x_{opt} by using a reference model such as $\dot{x}_m(t) = (-I + 2aP)x_m(t) + b$. However, while $x_m(t)$ converges to x_{opt} (see Proposition III.1 with $P = P^T$), the reference model is not implementable as the network matrix aP is unknown.

To overcome these challenges, we propose the *adaptive feedback* intervention protocol

$$u(t) = K(t)x(t) \tag{11}$$

with an adaptive gain matrix $K(t)$ determined by the following extended dynamics:

$$\dot{z}(t) = -z(t) + K(t)x(t) + b + u(t), \tag{12a}$$
$$\dot{w}(t) = -w(t) + e(t)x^\top(t)x(t), \tag{12b}$$
$$\dot{K}(t) = e(t)x^\top(t), \tag{12c}$$

where

$$e(t) := x(t) - z(t) - w(t).$$

Note that the intervention only uses information on b, and no knowledge on neither aP nor x_{opt} is required. The first dynamics (12a) aims to replicate the pseudo-gradient dynamics (9) and generate $z(t)$ such that it tracks the action profile $x(t)$. The second dynamics (12b) is included for technical reasons and is needed to guarantee boundedness of all solutions. This is formally stated in the following lemma.[3]

Lemma III.3. *Consider the pseudo-gradient dynamics* (9) *and let $P = P^\top$. Then, under the adaptive feedback intervention given by* (11) *and* (12)*, all solutions of the closed-loop system are bounded.*

[3]The proof of this technical result is quite elaborate and is dropped due to space constraints. The complete proof will be reported in an extended version of this work.

The next result establishes convergence to the social optimum x_{opt}.

Theorem III.4. *Let $P = P^\top$ and consider the pseudo-gradient dynamics* (9) *interconnected with the adaptive feedback intervention* (11), (12)*. Then, the action profile $x(t)$ asymptotically converges to the social optimum x_{opt} given by* (8).

Proof. Let $\xi := (x, e, \Psi)$ with $\Psi = K - aP$. Then, bearing in mind (9), (11) and (12), ξ admits the following dynamics

$$\dot{x} = (-I + 2aP)x + b + \Psi x, \tag{13a}$$
$$\dot{e} = -e - \Psi x - ex^\top x, \tag{13b}$$
$$\dot{\Psi} = ex^\top. \tag{13c}$$

We proceed by following similar arguments as in the proof of the LaSalle's invariance principle [27, Thm. 4.4], but the proof is tailored for a single (yet arbitrary) trajectory. Let $\xi_0 := (x_0, e_0, \Psi_0)$ with some $x_0, e_0 \in \mathbb{R}^n$ and $\Psi_0 \in \mathbb{R}^{n\times n}$, and $\xi(t)$ be a solution starting from the initial condition $\xi(0) = \xi_0$. It follows from Lemma III.3 that this solution is bounded. Thus, there exists a compact set $\mathcal{D}$ such that $\xi(t) \in \mathcal{D}$ for all $t \geq 0$. It also follows from [27, Lem. 4.1] that the positive limit set Ω of $\xi(t)$ is nonempty, compact, and invariant. Moreover, $\xi(t)$ approaches Ω as t tends to infinity.

We now consider the function

$$V(\xi) := \frac{1}{2}\|e\|^2 + \frac{1}{2}\|\Psi\|_{\text{F}}^2,$$

where we recall that $\|\Psi\|_{\text{F}}$ is the Frobenius norm. The derivative of V along the solutions of (13) is

$$\begin{aligned}\dot{V} &= -\|e\|^2 - e^\top\Psi x - \|e\|^2\|x\|^2 + \operatorname{Tr}(\Psi^\top ex^\top) \\ &= -\|e\|^2 - \|e\|^2\|x\|^2,\end{aligned} \tag{14}$$

where the last equality is obtained using $e^\top\Psi x = \operatorname{Tr}(\Psi^\top ex^\top)$. Therefore, we have $V \geq 0$ and $\dot{V} \leq 0$ which implies that $V(\xi(t))$ has a limit $V_\infty \geq 0$ as $t \to \infty$. Pick any point $\xi' \in \Omega$, then there is a sequence $\{t_n\}$, with $t_n \to \infty$ as $n \to \infty$, such that $\xi(t_n) \to \xi'$ as $n \to \infty$. We obtain from continuity of V that $V(\xi') = \lim_{n\to\infty} V(\xi(t_n)) = V_\infty$. Therefore, since ξ' is chosen arbitrary, we deduce that $V(\xi) = V_\infty$ for all $\xi \in \Omega$, which means that on the invariant set Ω, the function V is constant. Moreover, we have $\dot{V}(\xi(t)) = 0$ for all $\xi(t) \in \Omega$. Let $E := \{\xi \in \mathcal{D} \mid \dot{V}(\xi) = 0\}$, then we have $\Omega \subset E$. Now let M be the largest invariant set inside E, subsequently we have the following relations

$$\Omega \subset M \subset E \subset \mathcal{D}.$$

Noting that $\xi(t)$ approaches Ω as $t \to \infty$, we obtain that $\xi(t)$ approaches M as $t \to \infty$.

The last step is to find the set M. Note from the definition of E and (14) that $E = \{\xi \in \mathcal{D} \mid e = \mathbf{0}\}$. Thus, on the

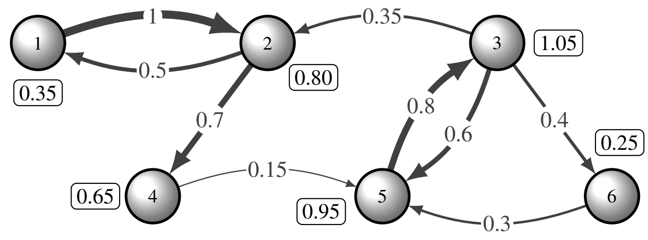

Fig. 1. The directed network.

invariant set M, the dynamics of (13) reads as

$$\begin{aligned}\dot{x} &= (-I + 2aP)x + b,\\ \mathbf{0} &= -\Psi x,\\ \dot{\Psi} &= \mathbf{0}.\end{aligned}$$

Noting that $-I + 2aP$ is Hurwitz as a consequence of Assumption II.2, the largest invariant set in E is given by

$$M = \left\{\xi \in \mathcal{D} \mid x = x_{\text{opt}},\ e = \mathbf{0},\ \Psi x_{\text{opt}} = \mathbf{0}\right\}.$$

Consequently, we conclude that $x(t)$ converges to x_{opt} as desired. □

IV. Illustrative Examples

In this section, we present two illustrative numerical examples and demonstrate the performance of our interventions.

A. Static feedback and dynamic interventions

We consider a population of $\mathcal{I} = \{1, \ldots, 6\}$ players that interact over the weighted directed graph shown in Fig. 1. The width of each link and its number indicate the weight of each link, and the numbers next to the nodes are their corresponding status-quo standalone marginal returns. The payoff function of each player is given by (2) with $a = -0.2$. The players choose their actions according to the dynamics (9) with random initial condition.

To steer the actions of the players to the social optimum, we use the static feedback and the dynamic interventions. For the former, we assume that the regulator knows aP to implement the intervention $u(t) = aP^\top x(t)$. In the latter given by $\dot{u}(t) = -(x(t) - x_s)$, we choose a random initial condition and set $x_s = x_{\text{opt}}$ by assuming that the regulator knows the value of the social optimum. Note that we have chosen $x_s = x_{\text{opt}}$ merely for the sake of illustration of the results, but convergence to any arbitrary choice of x_s is guaranteed (cf. Proposition III.2). The players' actions under the static feedback and the dynamic interventions are illustrated in Fig. 2. By implementing these mechanisms, the regulator can steer the action profile to the social optimum as demonstrated in Fig. 3.

B. Adaptive intervention

In this example, we illustrate convergence to the social optimum under the adaptive intervention (11), (12). We consider a similar network game to the previous example

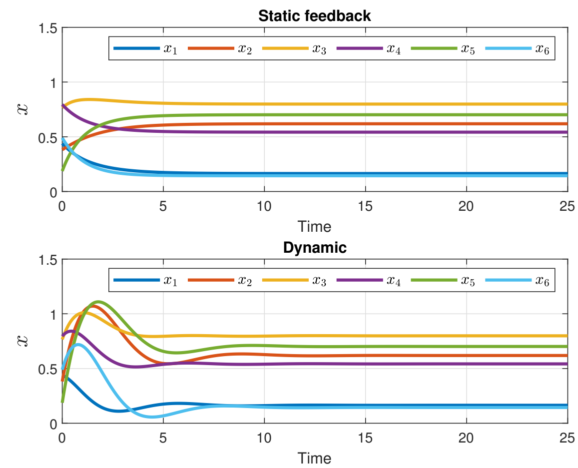

Fig. 2. Action profile under static feedback and dynamic interventions.

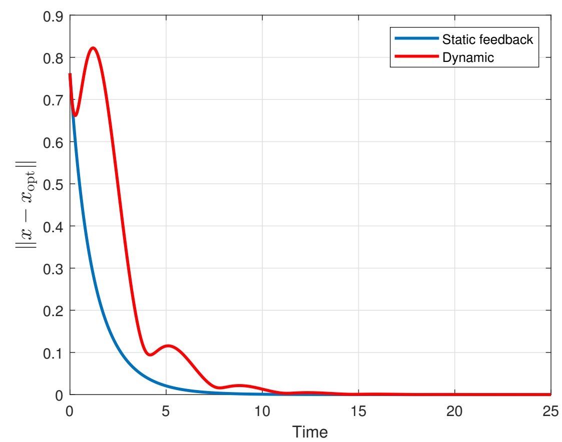

Fig. 3. Distance of action profile to social optimum under static feedback and dynamic interventions.

with $a = -0.2$, but the underlying network of the game is undirected here, as shown in Fig. 4. We randomly choose the initial condition of the overall system, made of the pseudo-gradient dynamics (9) and the adaptive intervention (11), (12). Fig. 5 depicts the resulting actions of the players. As shown in Fig. 6, the actions converge to the desired social optimum under the proposed adaptive intervention.

V. Conclusions

We considered the problem of steering the outcome of noncooperative quadratic network games to the social optimum. In our setup, the players use pseudo-gradient dynamics to choose their actions and maximize their individual payoffs. On the other hand, the regulator is aimed at choosing suitable interventions and regulating the actions of the players to the social optimum. We analyzed multiple scenarios where the regulator has access to limited information on the game and provided different intervention protocols that solve the problem.

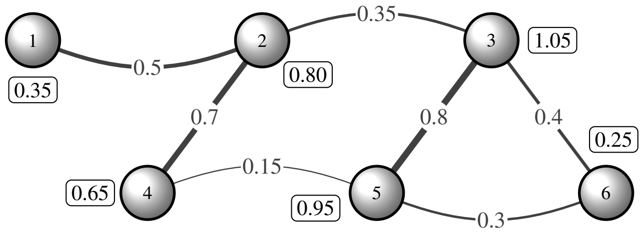

Fig. 4. The undirected network.

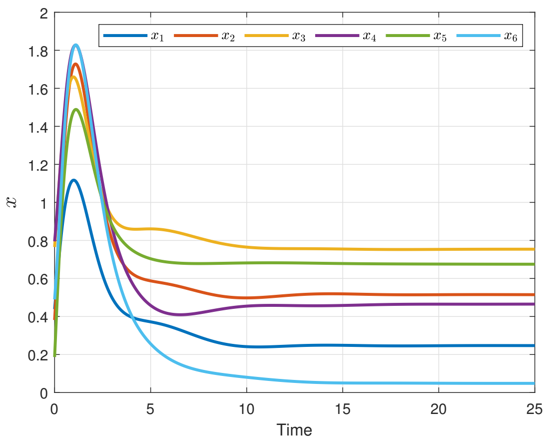

Fig. 5. Action profile under adaptive intervention.

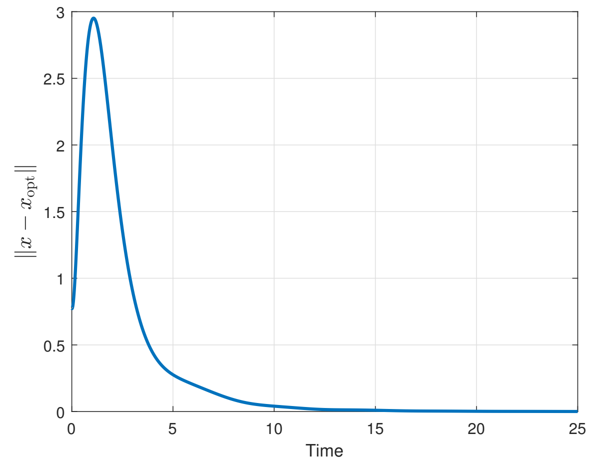

Fig. 6. Distance of action profile to social optimum under adaptive intervention.

Future research will explore extending the results to network games with general payoff functions, as well as including budget constraints on the interventions. Other research questions include interventions that are applied only to a subset of players as well as social welfare maximization for players with constrained actions.

References

[1] C. Ballester, A. Calvó-Armengol, and Y. Zenou, "Who's who in networks. wanted: The key player," *Econometrica*, vol. 74, no. 5, pp. 1403–1417, 2006.

[2] O. Candogan, K. Bimpikis, and A. Ozdaglar, "Optimal pricing in networks with externalities," *Operations Research*, vol. 60, no. 4, pp. 883–905, 2012.

[3] F. Bloch and N. Quérou, "Pricing in social networks," *Games and economic behavior*, vol. 80, pp. 243–261, 2013.

[4] Y. Bramoullé and R. Kranton, "Public goods in networks," *Journal of Economic theory*, vol. 135, no. 1, pp. 478–494, 2007.

[5] N. Allouch, "On the private provision of public goods on networks," *Journal of Economic Theory*, vol. 157, pp. 527–552, 2015.

[6] D. Acemoglu, A. Ozdaglar, and A. Tahbaz-Salehi, "Networks, shocks, and systemic risk," National Bureau of Economic Research, Tech. Rep., 2015.

[7] E. Altman, T. Boulogne, R. El-Azouzi, T. Jiménez, and L. Wynter, "A survey on networking games in telecommunications," *Computers & Operations Research*, vol. 33, no. 2, pp. 286–311, 2006.

[8] A. Galeotti, B. Golub, and S. Goyal, "Targeting interventions in networks," *Econometrica*, vol. 88, no. 6, pp. 2445–2471, 2020.

[9] G. Demange, "Optimal targeting strategies in a network under complementarities," *Games and Economic Behavior*, vol. 105, pp. 84–103, 2017.

[10] E. Koutsoupias and C. Papadimitriou, "Worst-case equilibria," in *Annual Symposium on Theoretical Aspects of Computer Science*. Springer, 1999, pp. 404–413.

[11] R. Johari, S. Mannor, and J. N. Tsitsiklis, "Efficiency loss in a network resource allocation game: the case of elastic supply," *IEEE Transactions on Automatic Control*, vol. 50, no. 11, pp. 1712–1724, 2005.

[12] G. Christodoulou and E. Koutsoupias, "The price of anarchy of finite congestion games," in *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, 2005, pp. 67–73.

[13] T. Roughgarden and É. Tardos, "How bad is selfish routing?" *Journal of the ACM (JACM)*, vol. 49, no. 2, pp. 236–259, 2002.

[14] G. Perakis and G. Roels, "The price of anarchy in supply chains: Quantifying the efficiency of price-only contracts," *Management Science*, vol. 53, no. 8, pp. 1249–1268, 2007.

[15] F. A. Hayek, "The use of knowledge in society," *The American economic review*, vol. 35, no. 4, pp. 519–530, 1945.

[16] W. Vickrey, "Counterspeculation, auctions, and competitive sealed tenders," *The Journal of finance*, vol. 16, no. 1, pp. 8–37, 1961.

[17] T. Roughgarden, "Algorithmic game theory," *Communications of the ACM*, vol. 53, no. 7, pp. 78–86, 2010.

[18] T. Alpcan and L. Pavel, "Nash equilibrium design and optimization," in *2009 International Conference on Game Theory for Networks*. IEEE, 2009, pp. 164–170.

[19] T. Alpcan, L. Pavel, and N. Stefanovic, "A control theoretic approach to noncooperative game design," in *Proceedings of the 48h IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*. IEEE, 2009, pp. 8575–8580.

[20] J. Barrera and A. Garcia, "Dynamic incentives for congestion control," *IEEE Transactions on Automatic Control*, vol. 60, no. 2, pp. 299–310, 2014.

[21] Y. Bramoullé, R. Kranton, and M. D'amours, "Strategic interaction and networks," *American Economic Review*, vol. 104, no. 3, pp. 898–930, 2014.

[22] J. Corbo, A. Calvó-Armengol, and D. C. Parkes, "The importance of network topology in local contribution games," in *International workshop on web and internet economics*. Springer, 2007, pp. 388–395.

[23] C. De Persis and S. Grammatico, "Distributed averaging integral Nash equilibrium seeking on networks," *Automatica*, vol. 110, p. 108548, 2019.

[24] D. Gadjov and L. Pavel, "A passivity-based approach to Nash equilibrium seeking over networks," *IEEE Transactions on Automatic Control*, vol. 64, no. 3, pp. 1077–1092, 2019.

[25] T. Basar and G. J. Olsder, *Dynamic noncooperative game theory*. Siam, 1999, vol. 23.

[26] D. Bertsekas, *Nonlinear Programming*, 2nd ed. Athena Scientific, 1999.

[27] H. Khalil, *Nonlinear Systems*, ser. Pearson Education. Prentice Hall, 2002.